

0 503 417 A2

EUROPEAN PATENT APPLICATION

② Application number: 92103531.7

Int. Cl.⁵: **G06F 11/10, G11B 20/18**

Ⓢ Date of filing: 02.03.92

① Priority: 11.03.91 US 668107

④ Date of publication of application:
16.09.92 Bulletin 92/38

③ Designated Contracting States:
DE FR GB

⑦ Applicant: FUJITSU LIMITED
1015, Kamikodanaka Nakahara-ku
Kawasaki-shi Kanagawa 211(JP)

② Inventor: Barlow, Howard L
1075 Hartford Drive

Boulder, Colorado 80303(US)

Inventor: Enochs, Kirk M.

1270 Chambers Drive

Boulder, Colorado 80303(US)

Inventor: Troicky, Alexander

2876 Humboldt Circle

Longmont, Colorado 80503(US)

74 Representative: Stebbing, Timothy Charles et al

Haseltine Lake & Co. Hazlitt House 28

Southampton Buildings Chancery Lane

London WC2A 1AT(GB)

④ Method and means for verification of write data.

52 A system for verifying write data includes a host computer (10), a controller (16) having a write data path network (18) and a read data path network (34), and a disk drive array (D0..., D3, DP) for storing write data. The host computer provides host cyclic redundancy check characters that are appended to blocks of input write data. The write data blocks are divided into subblocks that are processed through separate channels and ports (PORT0,...PORT3). Port

cyclic redundant check characters are appended to each data subblock, including a parity subblock. The read data path network (34) is coupled to receive the data subblocks with the appended host and port cyclic redundancy check characters from the write data path network means and has error detection circuitry (36, 38, 40, 42) for verifying the write data during the write mode.



This invention relates to the verification of write data in a storage means, for example in a fault tolerant disk drive array.

Presently known systems for recording binary data typically include a host computer for processing data, one or more disk drives and a controller interfacing between the computer and the disk drives. With some systems that process relatively large amounts of data, a plurality of drives are configured in an array for recording and storing data. A disk array generally consists of a number of disk drives and one or more redundant drives. Fault tolerant disk arrays are used that allow failure of one or more disk drives in an array without loss of data. Such types of disk arrays are used when data availability is critical. In one method, fault tolerance is achieved by calculating simple parity for the system data and storing it on the redundant drive(s). The reconstruction of data for a bad block requires the existence of parity data for that block, knowledge in the controller of the location of the bad block, and no more than one bad data block per parity block.

It is well known that the most common failure events which occur with the use of fault tolerant disk arrays are recoverable. Since fault tolerant disk arrays are employed for critical data applications, less common failure events affecting data integrity assume a much higher importance as compared to nonfault tolerant disk systems. Thus controller failures, even though much less likely than drive failures, may lead to data loss that is unacceptable. Although controller failure during a read operation leads to the lack of availability of system data, there is no permanent loss of the data. The faulty controller can be replaced and the data recovered. On the other hand, a controller failure during a write operation, if undetected, can cause a corruption of the system data on the drives. Such corruption of data would not be detected until the bad block is subsequently read. If there is a significant time lag between the write and read operations, additional bad blocks could be written in the interim. In a worst case situation, a controller failure during a write operation may never be detected in the controller. For example, if the controller scrambles data when writing to the disks and parity data is generated for the scrambled data, an undetectable and uncorrectable write failure would occur.

Typical controller architecture has two exposures to write data loss events, namely (1) an undetected failure of the controller data path function; and (2) an undetected failure in one of the drive data paths. Even though the drive data paths are fault tolerant in a disk array, a bad block must be detectable in order for it to be correctable. Typical disk controllers protect against write data

loss events by reading and checking the data just written. This form of "write verification" in disk controllers requires an additional revolution of the disks in order to perform the read operation, which constitutes a significant performance penalty and therefore is not generally attempted. Furthermore, guaranteed detection of a write failure in a drive data path is difficult to implement and rarely achieved in disk array systems.

One well known error detection approach uses cyclic redundancy check (CRC) characters for detection of data path errors. Presently known disk systems require CRC generation circuitry and separate additional independent circuitry for detecting data errors. This duplication of circuitry is relatively expensive. It is highly desirable to provide integrity checking in the controller during the write mode using simplified circuitry.

Accordingly, it is desirable to achieve verification of write data for fault tolerant disk arrays without the performance penalty incurred in typical disk controllers due to reading the previously written data.

It is also desirable to provide a verification system that detects loss of write data or failure in the controller data path during a write operation.

Furthermore, it is desirable to provide a verification system including a disk array that detects loss of write data or failure during the write mode without duplication of circuitry in the write and read data circuits.

According to a first aspect of the present invention there is provided a system for verifying write data to be recorded in a storage means, comprising a host interface for providing input blocks of write data; a means for generating host cyclic redundancy check characters appended to said write data blocks; and a controller coupled to said interface, said controller including a write data path network and a read data path network coupled to the output of said write data path network; said write data path network including means for dividing said data blocks into subblocks to be passed through respective ports, and port generator for appending port cyclic redundancy check characters to said data subblocks; said read data path network including first means for checking said host cyclic redundant check characters and second means for checking said port cyclic redundancy check characters.

According to a second aspect of the present invention, there is provided a system for verifying write data to be recorded in a storage means, comprising host interface means for providing input blocks of write data; a controller coupled to said interface, said controller including a write data path network and a read data path network coupled to the output of said write data path network; said

write data path network comprising means for dividing said data blocks into data subblocks; independent ports coupled to said dividing means for receiving said data subblocks; said read data path network comprising a data sequencer for sequentially channelling said data subblocks to recreate input blocks of write data; and means for checking said input blocks of write data; and means responsive to said checking means for flagging errors in the write data and initiating a retry operation if a failure is detected.

According to a third aspect of the present invention there is provided a method for verifying write data to be recorded in a storage means, comprising the steps of providing input blocks of write data; generating host cyclic redundancy check characters appended to said write data blocks; dividing said data blocks into subblocks to be passed through respective ports; appending port cyclic redundancy check characters to said data subblocks; and checking said host cyclic redundant check characters and said port cyclic redundancy check characters.

In another form of the present invention, there is provided a circuit for verifying write data for a fault tolerant data storage system comprising host interface means for supplying a block of nonfault tolerant write data and for generating cyclic redundant check characters to be appended to said write data block; write data path means coupled to said generating means, having means for dividing said data block into a number of port data subblocks for processing said data subblocks in a like plurality of ports, and means for appending port cyclic redundancy check characters to said data subblocks; means for providing at least one redundant data port; and read data path means coupled to said write data path means for calculating said port cyclic redundancy check characters independently for each port.

Thus, verification of write data is achieved by detection of an error in the controller data path (CDP) during a write operation in a fault tolerant storage means such as a disk array system. The CDP is separated into a write data path (WDP) and a read data path (RDP), which operate independently of each other. The RDP is enabled during a write operation to read output data from the WDP and to implement detection of data errors in the controller. If no error is detected in the write data exiting the WDP, we are assured of being able to subsequently read it because of the fault tolerance in disk drive data paths.

In operation, cyclic redundancy check (HOST CRC) characters are generated and added to the host write data. The HOST CRC characters are used in two ways. First, they are regenerated and checked in the RDP during write operations as an

overall check of the WDP processing of the host data block. Second, the HOST CRC characters are stored on the drives to be used on subsequent read operations as an overall check of the RDP as it rebuilds the host data block. The combined host data block and HOST CRC is divided into multiple data blocks by the WDP, one block per drive. The controller then generates additional cyclic redundancy check (PORT CRC) characters which are appended to each block and stored on the drives. The PORT CRC characters are regenerated during read operations and compared to the stored characters. If a disk drive data path has failed during a write or read operation, the PORT CRC check can be used as a pointer to the failed data path. The read data can then be reconstructed from the other good disk drives in a well known manner. This feature provides fault tolerance on read operations. The blocks having write data with HOST and PORT CRC characters are fed to the ports of a plurality of disk drives, which includes a parity drive for processing redundancy data. Simultaneously, error detection circuitry in the RDP checks the host data, redundant data, and the HOST CRC and PORT CRC characters to detect errors in the write data from the WDP.

In a specific implementation of RDP error detection circuitry, there is no need for duplication of PORT CRC generating circuitry in the RDP thus simplifying the write verification system and reducing cost of manufacture and maintenance expense. Instead, the individual data blocks are checked by summing the contribution of each port. The resulting sum is checked and if a failure is indicated, the write operation is retried.

Reference will now be made, by way of example, to the accompanying drawings, in which:

Figure 1 is a schematic circuit block diagram of a write data verification system, illustrating independent write data and read data paths, embodying the present invention;

Figure 2 is a schematic circuit diagram of a parity accumulator, such as used in the write data path of the system of Fig. 1;

Figure 3 is a schematic circuit diagram of error detection circuitry useful for verifying write data, including an exclusive OR summer and zero check circuit, such as employed in the read data path of the system of Fig. 1;

Figure 4 is a schematic circuit diagram of the CRC generator such as used in the WDP and RDP of the system of Fig. 1; and

Figure 5 is an example of a host data block transfer as the system in Fig. 1 sends it to a disk array.

With reference to Fig. 1, a write verification system comprises a host computer 10 that processes binary data, and a controller 16 including a

write data path (WDP) network 18 and a read data path (RDP) network 34. The WDP network 18 and RDP network 34, although connected, function independently to process write data and read data respectively. The RDP network 34 is enabled during the write mode to read output data from the WDP network 18.

WRITE DATA PATH DESCRIPTION

Host computer 10 provides blocks of host data to be passed to the controller 16 and to an array of disk drives D0-D3, DP for recording of the data. The computer 10 includes a generator 12 that generates a number of bytes of host cyclic redundancy check (HOST CRC) characters which are appended to each block of write data received at input terminal 14. The HOST CRC generator 12 does not need to reside in the host computer, but could just as easily reside in the controller 16 and equivalent results would be achieved.

The HOST CRC generator 12 is shown in more detail in Figure 4. As host data is received on input 76 and passed to output 78, the exclusive OR tree 70 and register 72 are operated in a well known manner such that the input data is divided by a generator polynomial and the remainder is computed. After a block of HOST data has passed through the generator 12, the multiplexer 74 is switched so as to append the remainder (which is now known as HOST CRC) onto the host data out of 78. This cycle is repeated for each block of host data.

As illustrated in Fig. 1, the write data block with the appended HOST CRC characters is directed to the WDP network 18. The block of input data from the host computer 10 and the HOST CRC characters are divided by a data sequencer 20 into a plurality of subblocks of port data, preferably of equal size. In this implementation, four data subblocks are passed respectively to PORT CRC generators 24a,b,c,d.

Concurrently the write input data from the host computer 10 is channeled to a parity accumulator 22 that calculates the redundancy data for the parity port, which is subsequently passed to PORT CRC generator 24p.

In accordance with this embodiment, PORT CRC characters are appended to the data and parity subblocks by PORT CRC generators 24a,b,c,d,p. The PORT CRC generators are implemented with the same type of circuit as the HOST CRC generator circuit shown in Figure 4. As each data byte is presented to one of the PORT CRC generators 24a,b,c,d,p for PORT CRC calculation, the byte is also passed through to the corresponding WDP holding register 26a,b,c,d,p. When each port data block has been completed, two PORT

CRC bytes are immediately transferred following the last port data byte to the respective WDP holding register 26. The data and parity subblocks are transferred from the WDP holding registers to the respective disk drives D0-D3,DP. The purpose of each holding register 26a,b,c,d,p is to provide temporary storage of data before the data can be sent to the respective disk drive.

During operation of the system, each of four consecutive host data bytes are summed by modulo 2 summation in the parity accumulator 22, and the resulting sum is fed to PORT CRC generator 24p. As an example, the first group of eight bytes of host write data may be designated as 01,02,03,04,05,06,07,08. The first byte 01 of 8 bits is transferred to PORT CRC generator 24a, WDP holding register 26a and, simultaneously to an exclusive OR (XOR) logic circuit 50 of the accumulator 22, as illustrated in Fig. 2. The byte value 01 is summed with the initial value 00 to produce a resultant 01. The second byte 02 is transferred to PORT CRC generator 24b and register 26b, and simultaneously summed with the current value 01 in accumulator 22 ($01 \text{ xor } 02 = 03$). Byte 03 is transferred to PORT CRC generator 24c, register 26c, and simultaneously summed ($03 \text{ xor } 03 = 00$). Byte 04 is transferred to PORT CRC generator 24d, register 26d and simultaneously summed ($00 \text{ xor } 04 = 04$). At this point the value of the byte in the parity accumulator 22 is transferred to PORT CRC generator 24p and register 26p. These five bytes, 01-04 and the parity byte, represent the first rank of data. A port rank corresponds to the set of bytes over which the parity byte is calculated, and the parity byte itself.

The accumulator 22 is reinitialized to 00 and the above sequence is repeated for every four bytes of input data. The transfer of data is continued by transfer of byte 05 to PORT CRC generator 24a, holding registers 26a, and simultaneous summing with the initialized 00 value ($00 \text{ xor } 05 = 05$). Then byte 06 is transferred to PORT CRC generator 24b, holding register 26b, and is simultaneously summed ($05 \text{ xor } 06 = 03$). Byte 07 is transferred to PORT CRC generator 24c, holding register 26c, and is simultaneously summed ($03 \text{ xor } 07 = 04$). Byte 08 is transferred to PORT CRC generator 24d, holding register 26d, and is simultaneously summed ($04 \text{ xor } 08 = 0C$). The parity byte 0C is then transferred to PORT CRC generator 24p, holding register 26p, and the accumulator is reinitialized to 00. The second rank therefore consists of the five bytes 05, 06, 07, 08, and the parity byte 0C.

This sequence is repeated until the entire host data block including HOST CRC characters has been transferred. The WDP Data Sequencer 20 divides the host data block including HOST CRC into fourths and the parity port data is calculated

from the host data. A system clock (not shown) is used to time the sequential transfer of the data bytes or subblocks.

After each host data block including HOST CRC characters has been processed in the above manner, the PORT CRC characters that are generated in PORT CRC generators 24a,b,c,d,p are passed respectively to WDP holding registers 26a-d,p. The data and parity subblocks with appended HOST CRC and PORT CRC characters are then fed through terminals 28a-d,p to PORT0-PORT3 and PORTP to respective data disk drives D0-D3 and parity disk drive Dp for storage.

READ DATA PATH DESCRIPTION

The data subblocks and redundant or parity subblock are fed back through terminals 30a-30d,p to the network 34 constituting the read data path (RDP). The port data and parity data are provided to holding registers 32a-d,p. Each holding register provides temporary storage of data. The stored data contents of the RDP holding registers 32a-d,p are passed to a 5 byte XOR summer 36 to produce a sum check across each rank, bytes 0-3,p for the entire data block. The output of the summer 36 is directed to a zero check circuit 38 which checks the result of the summing operation to be equal to 0 for each rank, bytes 0-3,p. If the sum is non-zero an error flag is set in a status register of the controller 16. The microprocessor 48 has access to the error flag. If the error flag is set after the write operation, the transfer is considered incorrect and the microprocessor initiates a retry. It should be noted that five 8-bit input data lines are connected between the outputs of the holding registers 32a-d,p and the inputs to xor summer 36, although Fig. 1 shows single lead lines for convenience and simplicity.

In Fig. 3, an example of detection circuitry comprising the xor summer 36 and zero check circuit 38 is illustrated. Five data lines, each carrying 8-bit data bytes, are applied to xor circuits 54, 56 and 58. Xor circuit 58 receives an 8-bit output from xor 56 as well as the 8-bit data byte from the holding register 32. The outputs of xor circuits 54 and 58 are supplied to xor circuit 60, which provides an 8-bit byte to the zero check circuit 38, having a zero detect circuit 62 and flip-flop 64. The output from the flip-flop 64 is an error flag which is set if the zero check circuit detects a non-zero result. The microprocessor 48 has access to the error flag. If the error flag 44 is set after the write operation, the transfer is considered incorrect and the microprocessor initiates a retry.

As shown in Fig. 1, the port subblocks provided to the RDP holding registers 32a-d are also provided to an RDP data sequencer 40 that se-

quentially channels the data subblocks to a HOST CRC check circuit 42 under control of the system clock. The data is reassembled in the RDP data sequencer 40 in the same sequence as it was provided by the host computer to the WDP. The sequence must be the same in order to regenerate the correct HOST CRC in the check circuit 42. The regenerated CRC value is compared to the HOST CRC characters. If the regenerated value is not the same as the HOST CRC, an error flag 46 is set in a status register of the controller 16. The microprocessor 48 has access to the error flag. If the error flag is set after the write operation, the transfer is considered incorrect and the microprocessor initiates a retry.

It should be noted that the HOST CRC check circuit 42 performs the same function as the HOST CRC generator 12 as illustrated in Figure 4 with the exception that the HOST CRC characters appended to the host data block are included in the checker's calculation. Also the result of the calculation is checked to be equal to zero at the end of the host data block.

A failure in the write data path causes the overall checks to fail in the read data path and the write operation is then immediately retried. If the RDP does not detect an error in the write data then the data and appended check characters are deemed to be fault tolerant. Fault tolerance allows any single drive path or drive to fail without loss of host data.

Since the host data block is checked by a HOST CRC, any error in the controller data path which occurs during the host data transfer will be detected. However since the PORT CRC characters are appended to the port data blocks after the HOST CRC, if an error occurs in the generation of the PORT CRC characters or when they are appended, the HOST CRC would not detect such errors during the write operation. An additional check is required to allow the detection of an error produced during port data block PORT CRC generation. One approach would be to duplicate the PORT CRC hardware in the RDP for each port, which would be relatively expensive and complex. For this example five check circuits similar to the circuit depicted in Fig. 4 would need to be duplicated in the RDP.

In an embodiment of this invention, a simple and efficient sum check of the PORT CRC characters is provided by using the linear properties of CRC polynomial operations. The sum of the individual data PORT CRC's and the parity PORT CRC will always be zero. In this implementation, the parity PORT CRC is generated from the parity port data and not from the sum of the individual data PORT CRC characters.

Each PORT CRC is the remainder of the di-

vision operation performed on each data subblock by the CRC polynomial defined as follows:

$$\text{PORT CRC} = r(x) = D(x) \text{ MOD } g(x)$$

where $r(x)$ = Remainder polynomial;

$D(x)$ = Data polynomial;

$g(x)$ = CRC generator polynomial

MOD = CRC MODULO 2 division

For this example there are four data polynomials divided by the CRC polynomial: $D0(x)$, $D1(x)$, $D2(x)$, $D3(x)$ one for each port data subblock. The remainder of the division operations are the PORT CRC's appended to the data polynomials.

$$\text{PORT CRC0} = r0(x) = D0(x) \text{ MOD } g(x)$$

$$\text{PORT CRC1} = r1(x) = D1(x) \text{ MOD } g(x)$$

$$\text{PORT CRC2} = r2(x) = D2(x) \text{ MOD } g(x)$$

$$\text{PORT CRC3} = r3(x) = D3(x) \text{ MOD } g(x)$$

The data polynomial of the parity port subblock ($Dp(x)$) is the MODULO-2 sum(xor) of the other ports' data polynomials:

$$Dp(x) = D0(x) + D1(x) + D2(x) + D3(x)$$

The remainder polynomial of the parity subblock is:

$$rp(x) = (D0(x) + D1(x) + D2(x) + D3(x)) \text{ MOD } g(x)$$

By superposition:

$$rp(x) = D0(x) \text{ MOD } g(x) + D1(x) \text{ MOD } g(x) + D2(x) \text{ MOD } g(x) + D3(x) \text{ MOD } g(x)$$

By substitution:

$$rp(x) = r0(x) + r1(x) + r2(x) + r3(x)$$

Or:

$$rp(x) + r0(x) + r1(x) + r2(x) + r3(x) = 0$$

This expression shows that the sum of the individual data PORT CRC's and the parity PORT CRC will be 0.

Fig. 5 illustrates an example of a 14 byte HOST DATA block with two bytes of HOST CRC appended, as follows:

HOST DATA: (01, 02, 03, 04, 05, 06, 07, 08, 09, 0A, 0B, 0C, 0D, 0E)

HOST CRC: (EF, FD)

The HOST DATA block is divided into four DATA subblocks. The parity subblock is generated from the HOST DATA. The parity subblock PORT CRC is calculated independently of the other PORT CRC characters.

The first rank of data consists of the HOST DATA block bytes 01, 02, 03, 04 and parity subblock byte 04. PORT0 data subblock consists of the host data block bytes 01, 05, 09, 0D, and PORT CRC bytes 22 and 6A as generated by the PORT CRC generator 24a. Each byte of the parity subblock is the xor sum across the corresponding data subblocks excluding the PORT CRC ranks 5 and 6. The parity PORT CRC as generated by 24p is calculated from the parity port data bytes 04, 0C, 04, 11. The sum across ranks 5 and 6 (PORT CRC ranks) should be zero:

$$22 \text{ xor } 14 \text{ xor } 84 \text{ xor } \text{FA} \text{ xor } 48 = 00, \text{ and}$$

$$6A \text{ xor } 5C \text{ xor } 0D \text{ xor } \text{AD} \text{ xor } 96 = 00.$$

Thus, an improved write verification system is provided wherein a sum check of the port CRCs takes advantage of the linear properties of CRC polynomials by individually summing the data port CRCs and the parity port CRC as a means to check the correct generation of port CRCs while simplifying the hardware design. The disclosed embodiment affords the checking of data blocks simultaneously with the recording and storage of data in a fault tolerant disk drive array.

The embodiment of the invention described herein is intended to be illustrative and it should be understood that the description is only one way of implementing the invention and is not meant to exclude others. For example, a byte for byte comparison could be performed between the WDP and RDP data as an alternative method for checking the write operation. In addition, while the invention is described with reference to a controller connected to disk drives, it is applicable to any controller connected in a fault tolerant way to any type of data processing and storage device.

Claims

1. A system for verifying write data to be recorded in a storage means, comprising:
 - a host interface for providing input blocks of write data;
 - a means for generating host cyclic redundancy check characters appended to said write data blocks; and
 - a controller coupled to said interface, said controller including a write data path network and a read data path network coupled to the output of said write data path network;
 - said write data path network including means for dividing said data blocks into subblocks to be passed through respective ports, and port generator for appending port cyclic redundancy check characters to said data subblocks;
 - said read data path network including first means for checking said host cyclic redundant

check characters and second means for checking said port cyclic redundancy check characters.

2. A system as claimed in claim 1, wherein said write data path network includes redundancy means coupled to said host interface for receiving said input data and for calculating redundancy subblocks to be passed through respective ports, and port generator means for appending port cyclic redundancy check characters to said redundancy subblocks. 5
3. A system as claimed in claim 1, wherein said write data path network comprises write data holding registers respectively coupled to said port generator means for temporarily storing respective subblock data including redundancy subblock data. 10
4. A system as claimed in claim 3, including terminal means comprising a number of terminals coupled to respective ones of said holding registers, and wherein said storage means is a fault tolerant disk drive array having a plurality of disk drives coupled to respective ones of said terminals. 15
5. A system as claimed in claim 4, wherein said plurality of disk drives includes at least one redundant disk drive for storing redundancy data, said redundant disk drive being coupled to a redundant port terminal means. 20
6. A system as claimed in claim 4 or 5, including a feedback path between said write data path network and said read data path network. 25
7. A system as claimed in claim 6, wherein said feedback path includes said terminal means, and said read data path network includes a number of read data holding registers coupled to said terminal means for temporarily storing said subblock data received from said terminal means; and a data sequencer for sequentially channelling said data subblocks including said redundant subblock stored in said holding registers. 30
8. A system as claimed in claim 7, wherein said first checking means comprising a host cyclic redundancy checking circuit that is coupled to said data sequencer for checking said host data cyclic redundancy check characters. 35
9. A system as claimed in any preceding claim, wherein said second checking means includes an exclusive OR summer coupled to said hold-

ing registers for summing said port cyclic redundancy check characters and a zero check circuit tied to said summer for checking said sum of said port cyclic redundancy check characters.

10. A system as claimed in claim 9, wherein said summer includes means for summing groups of consecutive host data bytes by modulo 2 summation. 40
11. A system for verifying write data to be recorded in a storage means, comprising:
 - host interface means for providing input blocks of write data; and
 - a controller coupled to said interface, said controller including a write data path network and a read data path network coupled to the output of said write data path network;
 - said write data path network comprising means for dividing said data blocks into data subblocks;
 - independent ports coupled to said dividing means for receiving said data subblocks;
 - said read data path network comprising:
 - a data sequencer for sequentially channelling said data subblocks to recreate input blocks of write data; and means for checking said input blocks of write data; and
 - means responsive to said checking means for flagging errors in the write data and initiating a retry operation if a failure is detected.
12. A system as claimed in claim 11, wherein said write data path network includes redundancy means coupled to said host interface for generating at least one redundant data subblock. 45
13. A system as claimed in claim 12, wherein said redundancy means is simple parity across the data subblocks, accumulated while dividing said write data blocks into subblocks. 50
14. A system as claimed in claim 11, 12 or 13, wherein said host interface means or said write data path network includes means for generating host cyclic redundancy check characters appended to said write data blocks; and said read data path network includes means for checking said host cyclic redundant check characters. 55
15. A system as claimed in claim 11, 12, 13 or 14, wherein said write data path network includes port generator means for appending port cyclic redundancy check characters to said data subblocks and, if provided, said redundancy subblocks; and said read data path network in-

cludes means for checking said port cyclic redundancy check characters.

16. A system as claimed in claim 15, wherein said port cyclic redundancy characters checking means includes duplicate CRC generating and comparing means. 5
17. A system as claimed in claim 15, wherein said port cyclic redundancy characters checking means includes an exclusive OR summer coupled to said write data path network for summing said port cyclic redundancy check characters and a zero check circuit tied to said summer for checking said sum of said port cyclic redundancy check characters. 10 15
18. A system as claimed in any of claims 11 to 17, wherein said independent ports are provided with a plurality of disk drives coupled to respective ones of said ports. 20
19. A method for verifying write data to be recorded in a storage means, comprising the steps of: 25
 - providing input blocks of write data;
 - generating host cyclic redundancy check characters appended to said write data blocks;
 - dividing said data blocks into subblocks to be passed through respective ports; 30
 - appending port cyclic redundancy check characters to said data subblocks; and
 - checking said host cyclic redundant check characters and said port cyclic redundancy check characters. 35
20. A method as claimed in claim 19, including the steps of receiving said input write data; calculating redundancy subblocks to be passed through respective ports; and appending port cyclic redundancy check characters to said redundancy subblocks. 40
21. A method as claimed in claim 19 or 20, including the step of temporarily storing subblock write data and redundancy subblock data. 45
22. A method as claimed in claim 19, 20 or 21, including the step of summing groups of consecutive host data bytes by modulo 2 summation. 50
23. A method as claimed in claim 19, 20, 21, or 22, including the step of flagging errors in the write data and initiating a retry if a failure is detected. 55

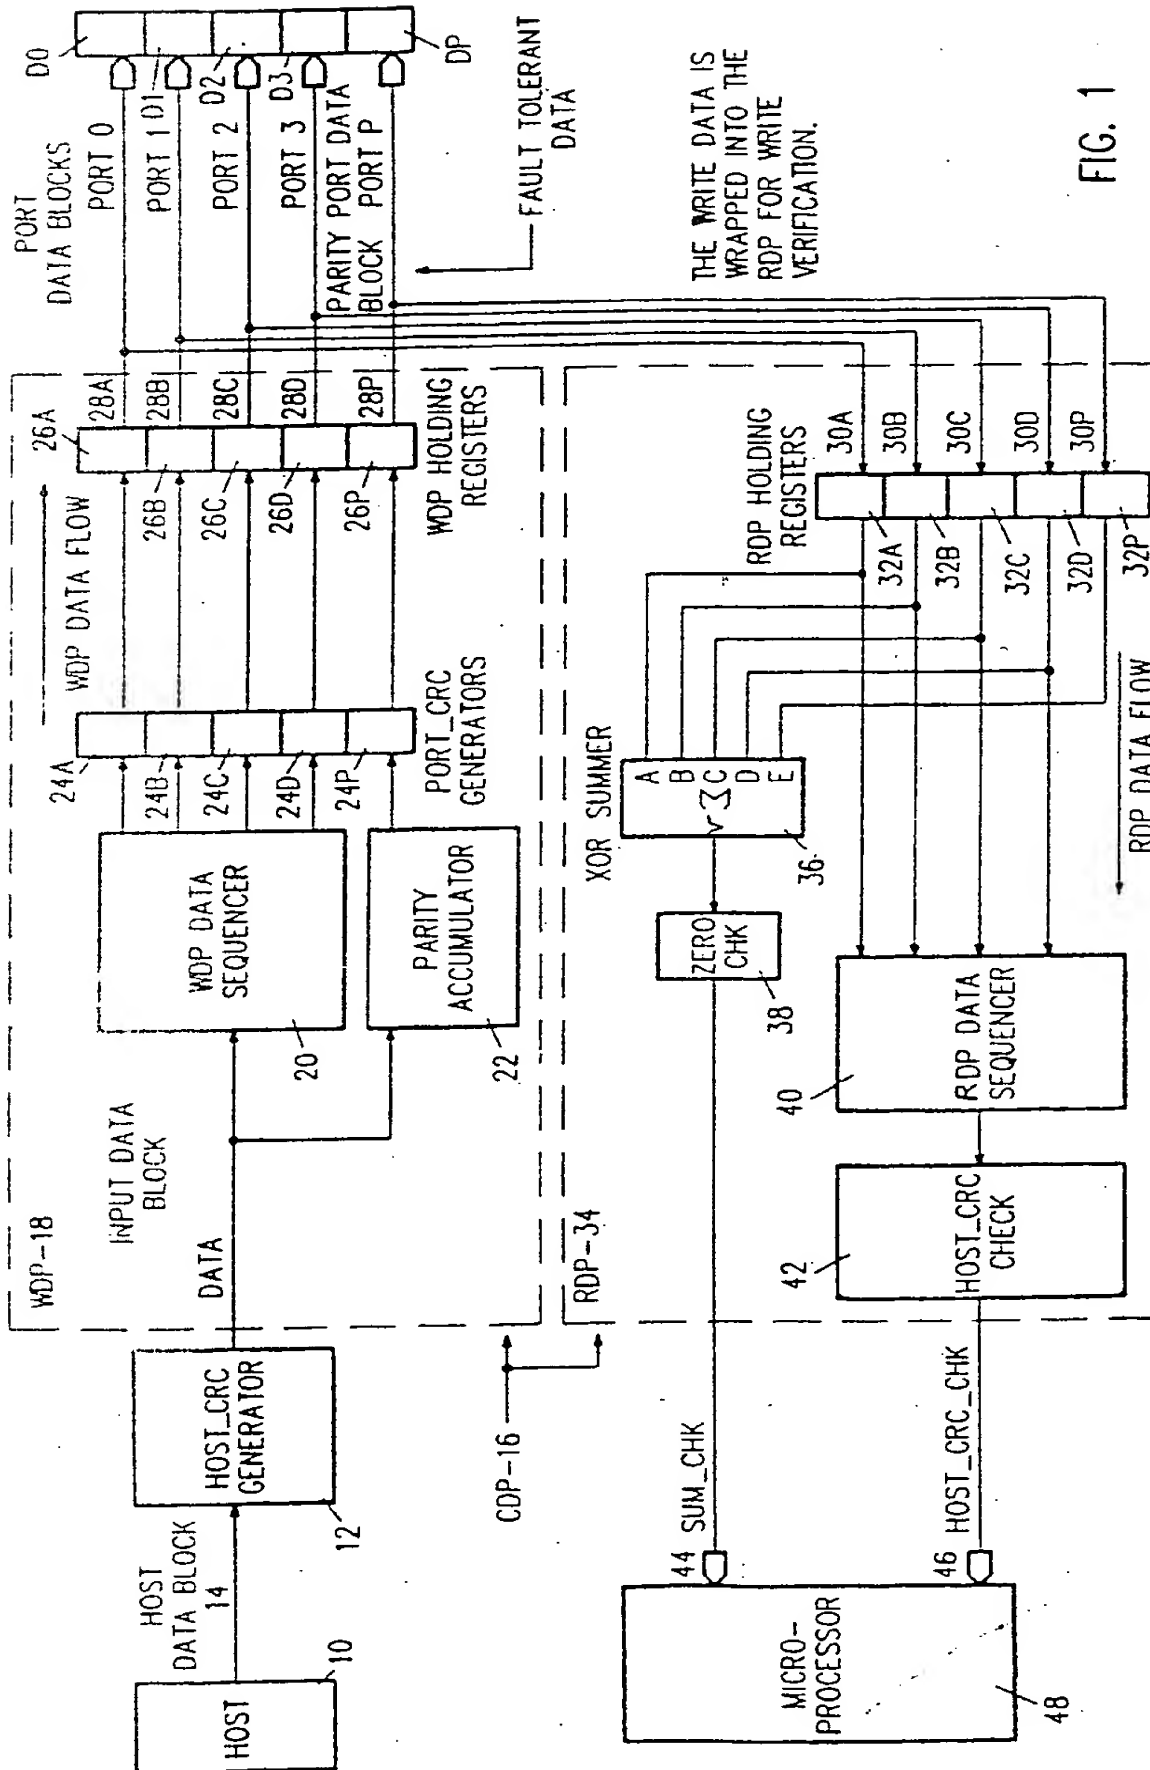


FIG. 1




FIG. 2




FIG. 3




FIG. 4

THE FOLLOWING HOST DATA BLOCK:

HOST DATA: (01, 02, 03, 04, 05, 06, 07,
08, 09, 0A, 0B, 0C, 0D, 0E)

HOST CRC: (EF, FD)

IS PROCESSED BY THE CDP AS FOLLOWS:

HOST DATA/HOST CRC					PORT CRC		
RANK =	1	2	3	4		5	6
D0(x) = (01	05	09	0D) MOD g(x) = r(x) =	22	6A
D1(x) = (02	06	0A	0E) MOD g(x) = r(x) =	14	5C
D2(x) = (03	07	0B	EF) MOD g(x) = r(x) =	84	0D
D3(x) = (04	08	0C	FD) MOD g(x) = r(x) =	FA	AD
Dp(x) = (04	0C	04	11) MOD g(x) = r(x) =	48	96
CHECKSUM =	00	00	00	00		00	00

NOTE: RANK 4; EF, FD, AND ALL CHARACTERS IN RANKS 5 AND 6 ARE CRC CHARACTERS.

FIG. 5



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 503 417 A3

EUROPEAN PATENT APPLICATION

(5) Application number: 92103531.7

(51) Int. Cl.⁵: G06F 11/10, G11B 20/18

(7) Date of filing: 02.03.92

(30) Priority: 11.03.91 US 668107

(53) Date of publication of application:
16.09.92 Bulletin 92/38

(84) Designated Contracting States:
DE FR GB

(38) Date of deferred publication of the search report:
06.07.94 Bulletin 94/27

(71) Applicant: FUJITSU LIMITED
1015, Kamikodanaka
Nakahara-ku
Kawasaki-shi Kanagawa 211(JP)

(72) Inventor: Barlow, Howard L.
1075 Hartford Drive
Boulder, Colorado 80303(US)
Inventor: Enochs, Kirk M.
1270 Chambers Drive
Boulder, Colorado 80303(US)
Inventor: Troicky, Alexander
2876 Humboldt Circle
Longmont, Colorado 80503(US)

(74) Representative: Stebbing, Timothy Charles et al
Haseltine Lake & Co.
Hazlitt House
28 Southampton Buildings
Chancery Lane
London WC2A 1AT (GB)

(54) Method and means for verification of write data.

(57) A system for verifying write data includes a host computer (10), a controller (16) having a write data path network (18) and a read data path network (34), and a disk drive array (D0..., D3, DP) for storing write data. The host computer provides host cyclic redundancy check characters that are appended to blocks of input write data. The write data blocks are divided into subblocks that are processed through separate channels and ports (PORT0....PORT3). Port

cyclic redundant check characters are appended to each data subblock, including a parity subblock. The read data path network (34) is coupled to receive the data subblocks with the appended host and port cyclic redundancy check characters from the write data path network means and has error detection circuitry (36, 38, 40, 42) for verifying the write data during the write mode.

EP 0 503 417 A3

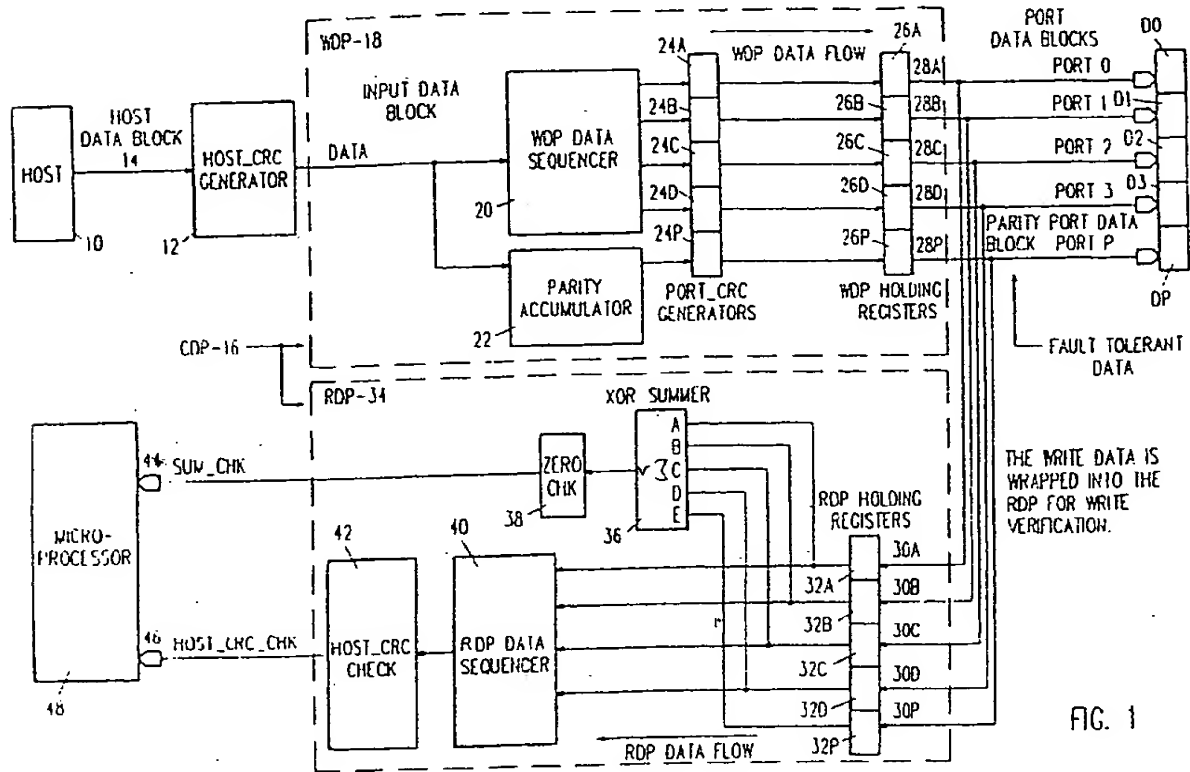


FIG. 1



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 92 10 3531

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. CL.5)
A	IBM TECHNICAL DISCLOSURE BULLETIN vol. 29, no. 4, September 1986, ARMONK, NY, USA pages 1678 - 79 'SELF-CHECKING ECC MODULE FOR A PROCESSOR MEMORY' * the whole document *	1, 11, 19	G06F11/10 G11B20/18
A	EP-A-0 294 287 (CAB-TEK, INC.) * page 3, line 1 - line 19; figures 4-8, 11 * page 7, line 2 - line 38 *	1, 11, 19	
A	WO-A-91 01524 (CRAY RESEARCH, INC.) * page 3, line 9 - line 21; figure 5 *	1, 11, 19	
A	IBM TECHNICAL DISCLOSURE BULLETIN vol. 32, no. 2, July 1989, ARMONK, NY, USA pages 321 - 323 XP000033443 'DISK FILE ERROR CHECKING' * the whole document *	1, 11, 19	
			TECHNICAL FIELDS SEARCHED (Int. CL.5)
			G06F G11B
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 4 May 1994	Examiner Absalom, R
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document			
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons * : member of the same patent family, corresponding document			

EPD FORM 1501 (03.91) (P01C01)